arXiv:1704.05646v1 [cs.LG] 19 Apr 2017

Effects of the optimisation of the margin distribution on generalisation in deep architectures

Lech Szymanski
University of Otago
lechszym@cs.otago.ac.nz

Brendan McCane
University of Otago
mccane@cs.otago.ac.nz

Wei Gao
Nanjing University
gaow@lamda.nju.edu.cn

Zhi-Hua Zhou
Nanjing University
zhouzh@lamda.nju.edu.cn

Abstract

Despite being so vital to success of Support Vector Machines, the principle of separating margin maximisation is not used in deep learning. We show that minimisation of margin variance and not maximisation of the margin is more suitable for improving generalisation in deep architectures. We propose the Halfway loss function that minimises the Normalised Margin Variance (NMV) at the output of a deep learning models and evaluate its performance against the Softmax Cross-Entropy loss on the MNIST, smallNORB and CIFAR-10 datasets.

1 Introduction

Support Vector Machines (SVM) guarantee best generalisation in a classification task for a chosen feature extraction function (Cortes & Vapnik, 1995; Vapnik, 1995). While the question of the choice of appropriate feature function (or its parameters) still remains, the training is guaranteed to give the optimal answer for the choice made. This assurance of generalisation comes from the principle of maximising the margin of separation.

Boosting methods build a feature space during the training process from an ensemble of weak classifiers (Schapire, 1990). It has been shown that their resistance to overfitting is due to the effect these methods have on the distribution of points around the margin (Schapire et al., 1998; Reyzin & Schapire, 2006; Wang et al., 2011). Gao & Zhou (2013) theoretically showed that AdaBoost is resistant to overfitting because it implicitly optimises the classification margin distribution by maximising average margin and minimising margin variance simultaneously. In particular, they emphasised that the minimisation of margin variance is very important, which was ignored by most previous studies on learning algorithm design. Zhang & Zhou (2013) proposed the LDM which maximises average margin and minimises margin variance simultaneously, and achieved consistently better performance than SVMs; later, Zhang & Zhou (2016) proposed Optimal Margin Machine (ODM) which demonstrates even better performance.

In this paper we take up the idea of margin distribution and apply it to deep learning. We theorise that in deep architectures with traditional backpropagation training, maximising the margin of separation is not likely to positively affect generalisation. However, we demonstrate that Halfway loss, which aims to minimise the normalised margin variance (NMV), does lead to improved generalisation in terms of outperforming the Softmax Cross-Entropy loss on the MNIST, smallNORB and CIFAR10 datasets.

2 Previous work

A number of different approaches have made an effective use of the principle of margin maximisation in artificial neural networks. Jayadeva et al. (2002) combined it with a decision tree-based training and Nishikawa & Abe (2002) incorporated it into the CARVE algorithm (Young & Downs, 1998). Both of these methods are based on a boosting-like training scheme, where the feature space is built one neuron (hypothesis) at a time, focusing on the remaining, incorrectly classified subset of the training data. Although in spirit these methods pertain to neural networks, the performance leverage they gain thanks to maximising the margin has probably more to do with the boosting aspects of the feature building rather than the deep nature of the neural network used.

The meticulously named Maximum Margin Gradient Descent with adaptive learning rate (MMGDX) algorithm proposed by (Ludwig & Nunes, 2010)'s works in a more traditional, fixed connectionist architecture trained with the backpropagation algorithm. The success of that algorithm most likely lies in the fact that the proposed Means Squared Error (MSE)-like loss, not unlike the Halfway loss introduced in this paper, might in fact be also minimising the distribution of margin variance. It also should be noted that MMGDX was tested only on single-hidden layer neural network with sigmoid activation function, and the superior performance only showcased on single-class problems.

The above methods train neural networks by maximising the geometric margin. Hence, in some sense, they consider the mean margin, but ignore the influence of margin variance (with the exception of MMGDX, which unintentionally might be reducing the variance). The Halfway loss minimises the margin variance and is not limited to the sigmoid activation function. We can test it on fully connected, as well as convolutional, neural networks with Rectifier Linear Unit (ReLU) activation function (Hahnloser et al., 2003; Glorot et al., 2011) and compare its performance to Softmax Central-Entropy loss on multi-class image recognition datasets.

3 Margin

Let's denote by $\mathcal{X} \in \mathbb{R}^d$ the instance space and by $\mathcal{Y} \in \{+1, -1\}$ the label set governed by some distribution $\mathcal{D}$ over $\mathcal{X} \times \mathcal{Y}$. Let's assume that we have a set $S = \{(\mathbf{x}_1, y_1), (\mathbf{x}_2, y_2), ..., (\mathbf{x}_m, y_m)\}$ of m points drawn identically and independently from D. Given some feature extraction function $\phi(\mathbf{x}) \in \mathbb{R}^k$ and a linear classifier parametrised by bias b and a unit vector $\mathbf{w} \in \mathbb{R}^k$, so that $|\mathbf{w}| = 1$, we can define the margin of instance $(\mathbf{x}_i, y_i)$, which is really a distance of the point from the classification boundary, as:

$$\gamma_i = y_i \left[\mathbf{w}^T \phi(\mathbf{x}_i) + b\right]. \tag{1}$$

Given classification error as

$$\epsilon(\gamma_i) = \begin{cases} 0 & \gamma_i > 0 \\ 1 & \text{otherwise,} \end{cases}$$

the goal of binary classification is to search for $\phi(\mathbf{x})$, the direction of unit vector $\mathbf{w}$, and value of b such that the expectation of $\epsilon(\gamma)$ over distribution $\mathcal{D}$, is minimised. Since typically $\mathcal{D}$ is not known, the best way to estimate the expectation is by computing the average of $\epsilon(\gamma_i)$ over the points from the training sample S and then try to maximise it. Breiman (1999) believed that AdaBoost also tried to maximise the minimum margin. Later, Reyzin & Schapire (2006) claimed that AdaBoost emphasises the average margin or median margin. The average margin is also called "margin mean", defined as:

$$\mu = \frac{1}{m} \sum_{i=1}^{m} \gamma_i. \tag{2}$$

3.1 Margin across different feature spaces

In SVMs the feature extracting function $\phi(\mathbf{x})$ is user-definable, but fixed during optimisation. The learning process for a given choice of $\phi(\mathbf{x})$ is the search for $\mathbf{w}$ and b that maximises the geometric margin while providing correct classification to the degree dictated by the choice of the soft margin parameter. A fixed feature space and margin maximisation guarantees an upper bound on probability of misclassification (Cristianini & Shawe-Taylor, 2000). The challenge with SVMs is to determine the best $\phi(\mathbf{x})$ by selecting the right kernel function and its parameters as well as an appropriate soft margin penalty factor for misclassification.

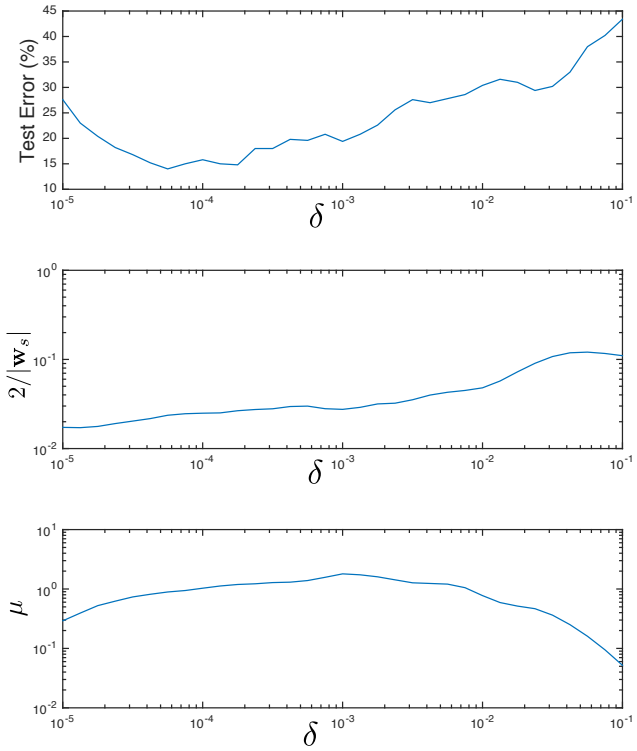

Figure 1: Test error and the corresponding values for the maximum geometric margin and the mean margin found by SVM training on a two-class subset from the smallNORB dataset using feature extraction function defined in Equation 3 over different values of parameter δ.

Given our intention to apply margin theory to deep learning, we are prompted to investigate the behaviour of the margin in an SVM while varying $\phi(\mathbf{x})$. Lanckriet et al. (2004) demonstrated that with some constraints and restrictions on $\phi(\mathbf{x})$, maximisation of the margin still provides an upper bound on probability of misclassification. However, in deep learning $\phi(\mathbf{x})$ can be a universal function approximator, hence we are interested in margin behaviour in general. Let us conduct a simple experiment.

Let's define the following feature extraction function

$$\phi(\mathbf{x}) = \left[e^{-\delta|\mathbf{x}-\mathbf{x}_1|} \quad e^{-\delta|\mathbf{x}-\mathbf{x}_2|} \dots e^{-\delta|\mathbf{x}-\mathbf{x}_m|}\right]^T, \tag{3}$$

where $\mathbf{x}_i$ for $i = 1...m$ correspond to the input data from the training set. The feature space in the above definition of $\phi(\mathbf{x})$ is the feature space of a Radial Basis Function (RBF) neural network parameterised by δ with the centres corresponding to all points in the training dataset. It is not as computationally efficient as the Gaussian kernel, but it gives a similar feature space while still allowing for the computation of $\phi(\mathbf{x})$ (which Gaussian kernel does not). After SVM training on $\phi(\mathbf{x})$ for a given value of δ, and with the soft margin parameter $C = 100$, we can compute $\mathbf{w}_s = |\sum_{i=1}^{m} y_i \alpha_i \phi(\mathbf{x}_i)|$, where $\alpha_i > 0$ are the support vectors found by the SVM. With the ability to compute $\mathbf{w}_s$, we can evaluate the value of the geometric margin, $\frac{2}{|\mathbf{w}_s|}$, as well as compute the unit vector $\mathbf{w} = \frac{\mathbf{w}_s}{|\mathbf{w}_s|}$ and thus the mean margin, μ, for different values of δ.

Figure 1 shows how the test error relates to the maximum geometric and mean margin values over a range of different δ's in $\phi(\mathbf{x})$ for a two-class subset problem from the smallNORB dataset (LeCun et al., 2004). Note that the value of maximum geometric margin steadily climbs with δ despite the fact that the test error dips, reaches a minimum, and then starts climbing as the value of δ increases. The *best* performance does not correspond to the largest value of the *best* geometric margin found across different $\phi(\mathbf{x})$s.

We need to acknowledge that the lack of correlation between best test data performance and maximum value of the margin in the experiment above does not mean that there isn't an upper bound on misclassification for changing $\phi(\mathbf{x})$. However, given that current proofs for existence of the bound require certain constraints on the structure of $\phi(\mathbf{x})$ (Lanckriet et al., 2004), the result of our simple experiment prompts us to hypothesise that in general it is the relative value of the margin within given $\phi(\mathbf{x})$ and not its absolute value across different realisations of $\phi(\mathbf{x})$ that needs to be maximised in order to improve generalisation. This would suggest that it might not be advantageous to maximise the margin across different realisations of $\phi(\mathbf{x})$.

Figure 1 also shows the mean margin value resulting from SVM training on different realisations of $\phi(\mathbf{x})$. Although it doesn't increase steadily with δ, its maximum value does not coincide with the lowest test error either. The value of the mean margin is not consistently related to best test performance for varying $\phi(\mathbf{x})$.

3.2 Margin in deep architectures

The simple experiment from the previous section suggests that maximising margin in architectures where $\phi(\mathbf{x})$ is not constant may not lead to a better generalisation. We can go even further and show that a simple linear transformation facilitated by $\phi(\mathbf{x})$ is sufficient to produce arbitrary margin value without changing the relative position of the points with respect to the separating line given by $\mathbf{w}$.

Lemma 3.1. *The mean margin of a set of points classified by unit vector* $\mathbf{w}$*, bias* βb*, and a feature extracting function* $\phi(\mathbf{x}) = \beta\widehat{\phi}(\mathbf{x})$*, such that* $\mu > 0$*, can be made arbitrarily large by varying the value of* $\beta > 1$*.*

Proof. The lemma is rather obvious, since

$$\beta\gamma_i = y_i \left[\mathbf{w}^T \left(\beta\widehat{\phi}(\mathbf{x})\right) + \beta b\right],$$

which produces a mean margin $\beta\mu > \mu$ if $\mu > 0$ and $\beta > 1$. Note that, while bias of the linear classifier is allowed to vary, $\mathbf{w}$ remains a unit vector, as stipulated in our definition of the margin. It is also important to note that this transformation does not change the sign of any γ_i - all the points are classified exactly the same as before and after multiplication by β. Thus, this transformation doesn't change anything about the classification decision in the space of $\phi(\mathbf{x})$. □

To understand the significance of Lemma 3.1, let us suppose that we are trying to maximise the mean margin in a computational model where the feature extracting function is defined by a neural network, such that

$$\phi(\mathbf{x}) = f\left(\widehat{\mathbf{W}}^T\widehat{\phi}(\mathbf{x}) + \widehat{\mathbf{b}}\right),$$

where $f(x)$ is a monotonically increasing activation function, $\widehat{\mathbf{W}}^T$ and $\widehat{\mathbf{b}}$ are the parameters of the penultimate layer, and $\widehat{\phi}(\mathbf{x})$ is the output due to all the previous layers of the network. A simple linear transformation within $\phi(\mathbf{x})$ is sufficient to increase the margin. The representation power of the penultimate layer is more than sufficient to provide this transformation, by changing $\widehat{\mathbf{W}}^T$ and $\widehat{\mathbf{b}}$, without any changes to $\widehat{\phi}(\mathbf{x})$ or the location of the separating hyperplane in the feature space of $\phi(\mathbf{x})$. Thus it's possible for the data in the feature space to stretch away from the separating hyperplane and give a larger margin without any meaningful change to the feature extraction or classification.

As a result of Lemma 3.1 we form a hypothesis that maximisation of geometric or mean margin is not a meaningful objective for improving generalisation in deep architectures.

4 Margin variance

Following the theory developed by (Gao & Zhou, 2013) and (Zhang & Zhou, 2016) we next consider the effect of minimising the variance of the margin in deep architectures. The variance of the margin is defined as

$$\sigma = \frac{1}{m}\sum_{i=1}^{m}\left(\gamma_i - \mu\right)^2. \tag{4}$$

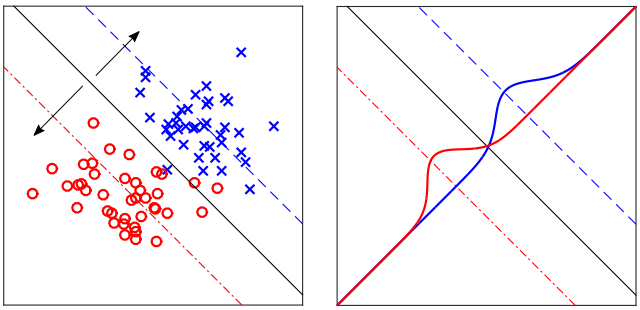

(a)

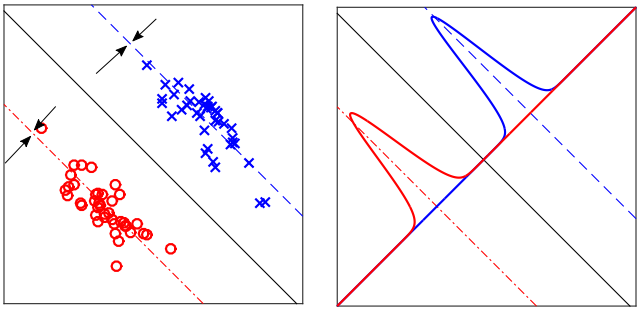

(b)

Figure 2: Illustration of transformations of simple 2-dimensional $\phi(\mathbf{x})$ that increases the mean margin a); reduces the margin variance b). The left-hand side plots shows the distribution of two-class set of points with different margin distribution and direction of the desired space transformation with the black arrows. The right-hand side show the distribution of the data around the margin. The separating line is shown as a black line with the corresponding margin denoted with the red and blue dashed-lines.

In order to increase the mean margin, as illustrated in Figure 2a, it is sufficient for the feature space $\phi(\mathbf{x})$ to change so that the points stretch away from the separating hyperplane defined by $\mathbf{w}$. This can be easily facilitated via a linear transformation, as stipulated in Lemma 3.1. Figure 2b depicts the type of transformation that $\phi(\mathbf{x})$ needs to undergo in order to reduce the margin variance. In addition to the stretch away from the separating hyperplane, the space must squash around two separate hyperplanes on the positive and negative margin. It is apparent that this is a somewhat less trivial non-linear transformation, and thus more likely to be conducive to meaningful changes of $\phi(\mathbf{x})$ with respect to generalisation.

4.1 Normalised margin

If our hypothesis, that the mean margin value is arbitrary for changing $\phi(\mathbf{x})$, is correct it stands to reason that the variance value might be arbitrary for different $\phi(\mathbf{x})$ as well. Indeed, if we repeat the experiment with SVM on the two-class subset of smallNORB and the feature extraction function defined in Equation 3, we can clearly see (in Figure 3) that the minimum margin variance does not exactly coincide with the minimum test error in terms of the δ value that specifies the curvature of $\phi(\mathbf{x})$. It should be noted that the SVM training does not strive to minimise the variance of the margin, but rather to maximise the geometric margin. However, given that margin is not consistent for different $\phi(x)$, it's not unreasonable to assume that variance won't be either. Hence we propose the normalised margin variance (NMV) defined as

$$\overline{\sigma} = \frac{1}{m} \sum_{i=1}^{m} \left(\overline{\gamma}_i - \overline{\mu} \right)^2 , \tag{5}$$

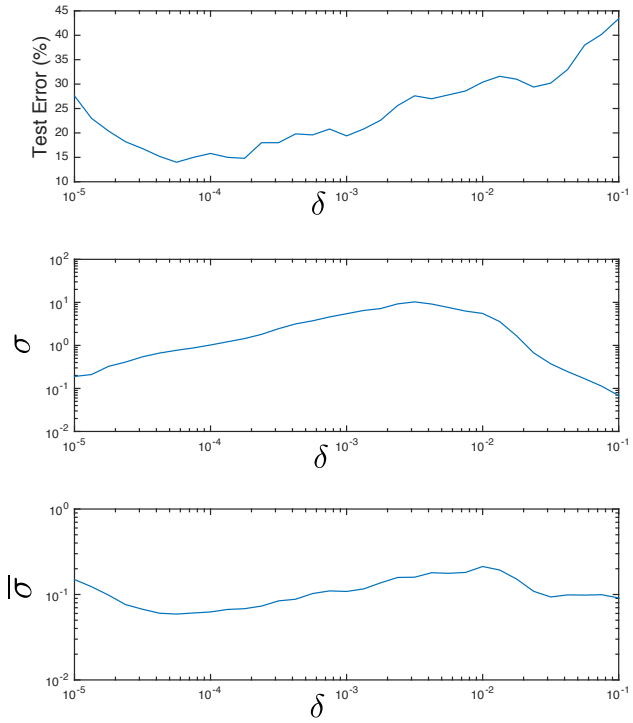

Figure 3: Test error and the corresponding values for the margin variance and the normalised margin variance after SVM training on a two-class subset from the smallNORB dataset using feature extraction function defined in Equation 3 over different values of parameter δ.

where

$$\overline{\gamma}_i = \frac{\gamma_i}{\max(\{|\gamma_1|, ..., |\gamma_m|\})},$$

and

$$\overline{\mu} = \frac{1}{m}\sum_{i=1}^{m}\overline{\gamma}_i.$$

Equation 5 has been designed to make the margin value $-1 \leq \overline{\gamma}_i \leq 1$. It is also worth to note that in a scenario where $\gamma_i = y_i\left[\mathbf{w}^T\beta\hat{\phi}(\mathbf{x}) + \beta b\right]$, the proposed normalisation removes the contribution of β to the margin value. The normalisation becomes

$$\overline{\gamma}_i = \frac{\gamma_i}{\gamma_j} = \frac{y_i\beta\left[\mathbf{w}^T\hat{\phi}(\mathbf{x}) + b\right]}{y_j\beta\left[\mathbf{w}^T\hat{\phi}(\mathbf{x}_j) + b\right]}, \tag{6}$$

where j is the index of the sample that produces maximum absolute value of the margin. The βs cancel out. This means that the linear transformation aspect of $\phi(\mathbf{x})$, which can give an arbitrary margin value at the output, is removed from the optimisation.

Figure 3 shows the normalised margin variance, $\overline{\sigma}$, for different values of δ in the SVM and two-class smallNORB experiment. The minimum of normalised margin variance does indeed fall close to the δ that gives the smallest test error.

4.2 Halfway loss function

In order to carry out empirical evaluation of the effect that minimising normalised variance has on generalisation in deep architectures, we propose the Halfway loss function defined as

$$J = \frac{1}{m}\sum_{i=1}^{m}\left(\overline{\sigma}_i - \frac{1}{2}\right)^2. \tag{7}$$

It is hard not to notice the resemblance of Equation 7 to the Mean Squared Error (MSE) loss function. MSE training does in fact strive to minimise the variance of the model's output around the value given by the target label. The point of difference between Halfway and MSE loss is the normalisation of the margin, which in effect is the same as normalisation of the model's output.

The motivation for normalisation, as discussed in the previous section, is to obtain consistency of the margin variance across different $\phi(\mathbf{x})$. However, a consequence of this normalisation is that optimisation does not enforce an absolute target value for the output, but rather a relative value with respect to other outputs. We hypothesise that part of the reason why Softmax is so successful in deep learning is that it allows the model to produce output in any range, as long as the relative value of the correct class neuron is larger than the value of other outputs. This allows the deep model to operate in its *natural* range, the values of the output being a result of the dynamics arising from the learner's architecture and the type of optimisation. This natural range might be also the reason why RELU activation function works so well with Softmax. Normalisation of the margin assures that Halfway loss, in contrast to MSE, does allow the model to operate in its natural range, though still drives the model to produce positive and negative output in correspondence to the sign of the target label.

The Halfway loss is basically a MSE loss that minimises the margin of a classifier around half way to the current maximum value of absolute margin, $\frac{1}{2}\max(\{|\gamma_1|, ..., |\gamma_m|\})$. The choice of $\frac{1}{2}$ for the target value for the normalised margin is based on the assumption that the mean of the margin is somewhere between 0 and the current maximum value.

4.3 Halfway loss for multi-class classification

For multi-class classification, where label $\mathbf{y}_i = \{-1, 1\}^K$ we propose a one-against-rest training scheme with a cost sensitive-learning-like (Elkan, 2001) multi-class weighting factor to correct the natural imbalance of the positive to negative label ratio. In an m-point dataset with even distribution of K classes, that is $\frac{m}{K}$ examples of each class, a given output will be trained on $\frac{m}{K}$ positive and $\frac{(K-1)m}{K}$ negative labels. This imbalance would mean that negative labels gain more variance reduction as opposed to the positive ones. In order to correct this, we propose the following Halfway loss for output k:

$$J_k = \frac{1}{m}\sum_{i=1}^{m}\overline{y}_{ki}\left(\overline{\gamma}_{ki} - \frac{1}{2}\right)^2, \tag{8}$$

where

$$\overline{y}_{ki} = \begin{cases} 1 & y_{ki} = 1 \\ \frac{1}{K-1} & y_{ki} = -1. \end{cases}$$

The symbols $\overline{\gamma}_{ki}$ and y_{ki} represent the normalised margin and the target label of the k^{th} output for input i. The multi-class weighting factor $\overline{y}_{ki}$ can be derived from the label as follows:

$$\overline{y}_i = \frac{(K-2)y_i + K}{2K-2}.$$

Note that the two-class Halfway loss defined in Equation 7 is analogous to $K = 2$ case of the multi-class Halfway loss defined in equation 8.

5 Empirical evaluation of Halfway loss

The three datasets that we will use for evaluation of the Halfway loss are the MNIST (LeCun et al., 1998), smallNORB (LeCun et al., 2004) and CIFAR-10 (Krizhevsky, 2009) datasets. Each set

Table 1: Averate test error (in % of misclassified samples) for the Softmax vs Halfway training over 10 trials on the MNIST, smallNORB and CIFAR-10 datasets and three deep architectures.

	Softmax	Halfway
MNIST		
FC-128-32	2.26 ± 0.09	2.32 ± 0.11
FC-500-500-2000	1.96 ± 0.11	1.62 ± 0.12
CNN	0.62 ± 0.05	0.55 ± 0.04
smallNORB		
FC-128-32	31.99 ± 2.75	31.75 ± 5.56
FC-500-500-2000	27.84 ± 3.55	22.64 ± 1.50
CNN	12.51 ± 1.07	11.01 ± 1.16
CIFAR10		
FC-128-32	50.12 ± 0.50	51.65 ± 0.69
FC-500-500-2000	51.29 ± 0.49	46.60 ± 0.43
CNN	30.06 ± 0.94	27.83 ± 0.47

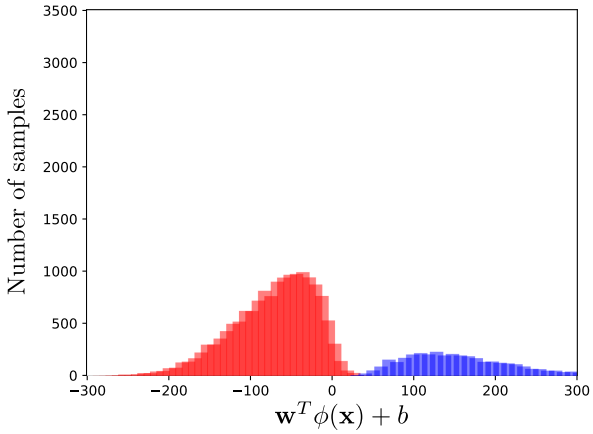

(a) Softmax loss

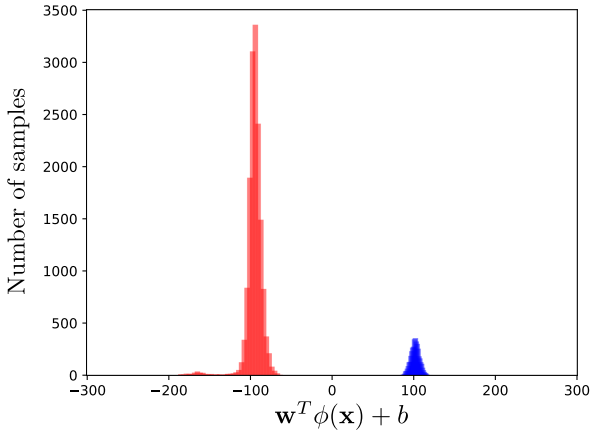

(b) NMV loss

Figure 4: Distribution of the positively (blue) and negatively (red) labelled output of the train data of a single output of a smallNORB-trained CNN.

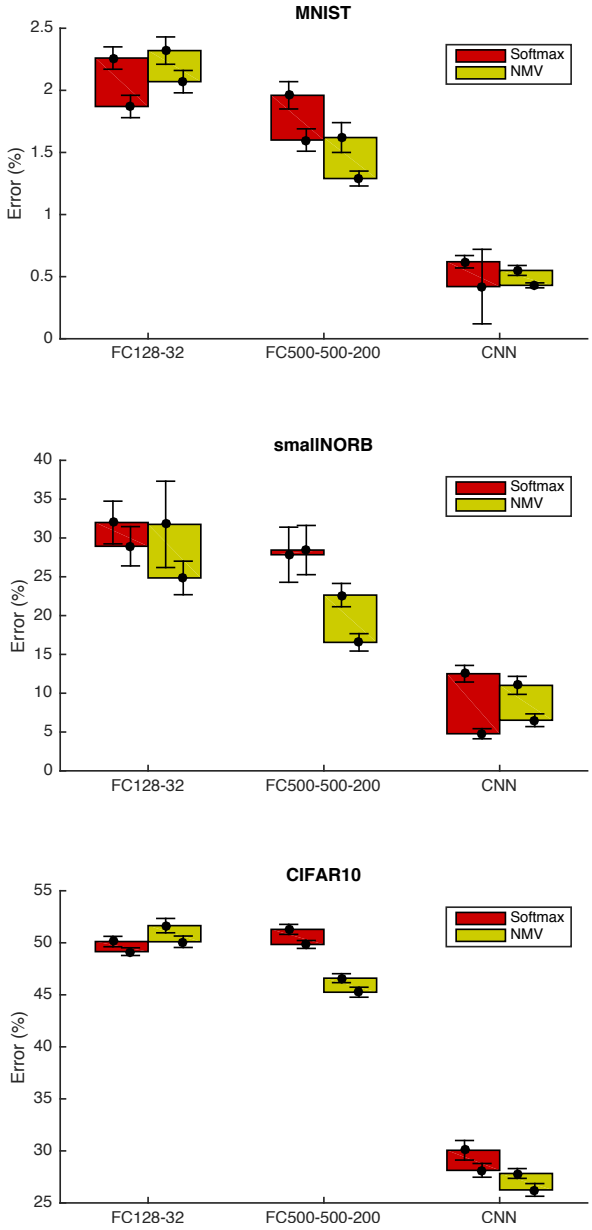

Figure 5: Generalisation error between the validation and test data represented as length of the bars. The top of the bar is placed at the test error mean and bottom at the validation error mean for 10 trials of different architectures and datasets. The standard deviation of the test error is shown on the left hand-side of each bar, the standard deviation for the validation error shown on the right-hand side of each bar. The errors for Softmax are shown in red and for Halfway in yellow.

comes pre-divided into a training (60000 MNIST, 24300 smallNORB, 50000 CIFAR-10) and testing (10000 MNIST, 243000 smallNORB, 10000 CIFAR-10) set of samples. For our evaluation, we have split each training part of the dataset into a set of images used to train the models (55000 MNIST, 19440 smallNORB, 45000 CIFAR-10), and a validation set (5000 MNIST, 4860 smallNORB, 5000 CIFAR-10) used to determine the best manifestation of the model. The choice of the validation sample for MNIST and CIFAR-10 was made randomly, whereas for the smallNORB dataset it was all the images of a specific instance of each of the class of toys, with a random choice of which instance was used for validation. For the smallNORB dataset only the left camera images were used.

Training was done with mini-batch optimisation, with a 500 sample batch-size for MNIST and CIFAR-10, and a 405 sample batch-size for smallNORB. The normalisation of the margin was carried out independently in each batch, which makes the Halfway loss somewhat similar to the batch normalisation transform proposed by Ioffe & Szegedy (2015). However, whereas the objective of batch normalisation requires computing the mean and variance across the batch sample in order to normalise its first and second order statistics, our normalisation of the margin only divides the data by the maximum absolute value of the output. Also, the aim of producing unity variance of the batch sample in batch normalisation, regardless of the label, is counter-objective to ours, which is to minimise the variance around the margin of different labels. Nevertheless, it is quite possible that Halfway loss minimisation with min-batch training shares some of the effects of reducing the *internal covariate shift* of batch normalisation in the output layer of the network.

All the evaluations were done using TensorFlow (Abadi et al., 2015), which provides automatic computation of the gradients required for optimisation. In our implementation there was no constraint placed on $\mathbf{w}$ to make it a unit vector, since margin normalisation, as shown in Equation 6, has the same effect making individual normalisation of $\mathbf{w}$ irrelevant.

Three different models were used for classification of each dataset. The small fully connected neural network (FC-128-32) consisted of 2-hidden layers with 128 and 32 neurons in the consecutive layers. The big fully connected neural network (FC-500-500-2000) consisted of 3-hidden layers with 500, 500 and 2000 neurons in the consecutive layers. Finally a CNN model was used for classification in each dataset. For the MNIST and smallNORB dataset a CNN consisted of two convolutional layers. The first convolutional layer had 32 filters of 5x5 size and stride of 1 followed by a 2x2 input max-pooling with stride 2; the second convolutional layer had 64 filters of 5x5 size and stride 1 followed by a 2x2 input max-pooling with stride of 2; this was followed with a 512-neuron fully connected layer and 0.5 probability dropout during training. For the CIFAR-10 model the CNN consisted of two convolutional layers also. The first convolutional layer had 54 filters of 5x5 size and stride 1 followed by a 3x3 input max-pooling with stride of 2 and local response normalisation; the second convolutional layer consisted of 64 filters of 5x5 size and stride of 1 followed by local response normalisation and a 3x3 input max-pool layer with stride of 2; this was followed by two fully connected layers of 384 and 192 neurons. The activation function used in all networks was ReLU.

The optimisation for all tests was done using Tensorflow's implementation of the Adam optimiser (Kingma & Ba, 2014). The learning rate for all runs was set to 0.001 and the maximum number of training epochs, with one epoch training over the entire set of mini-batch blocks, was set to 2000. There was no regularisation in any of the models, as the purpose of this evaluation was to compare the sole effect of the compared loss functions. Hence, rather than going for the state of the art results, we are aiming for a comparative study of the effect of minimisation of margin variance as compared to Softmax with Cross-Entropy loss function training. The $\mathbf{y} \in \{0, 1\}^K$ coding was used for target labels during Softmax training and $\mathbf{y} \in \{-1, 1\}^K$ coding was used for target labels during Halway loss training.

Table 1 reports the average test error and variance over 10 trials with the same initial values of weights and biases for a given trial between the Softmax and Halfway optimisation. Test error was measured by taking the output of with the maximum value to indicate the index of the identified class. The reported test error comes from the model state at the training epoch that produced the lowest validation error. The Halfway loss consistently leads to lower test error, sometimes by quite a significant amount.

Figure 4 shows the distribution of the values across a single output, $\mathbf{w}\phi(\mathbf{x}) + b$, from the entire train set on the smallNORB-trained CNN using Softmax and Halfway optimisation. Halfway trained model does indeed produce output with a smaller variance around the margins while maintaining the values in a similar range to the Softmax trained model.

It is also interesting to examine the difference between the validation and test error. In some way, it gives an idea of the generalisation error. Validation error stands for the empirical risk, since it was used during training to choose the best model (deemed to be the one that gives minimum validation error). The test error, although still just an average, simulates the true risk, since it has not been seen by the learner during the training. Figure 5 shows a bar plot comparing the generalisation error between Softmax and Halfway training for tested datasets and architectures. The length of the bars in the plot corresponds to the generalisation error; top and the bottom positions of each bar demarks the

test and validation error respectively. The desired characteristic is for the top of the bar to be as lows as possible (low test error) and the bar to be as short as possible (validation error being close to the test error). Although not in every single case, Halfway looks to outperform Softmax in a combination of lower test error and/or smaller generalisation error.

6 Discussion

When it comes to the bigger models, FC-500-500-2000 and CNN, those trained with the Halfway loss consistently outperform those trained with the Softmax Cross-Entropy in terms of the mean and also the standard deviation of the test error over multiple trials of different initial conditions. At the same time, for the small network, FC-128-32, Halfway training performs consistently worse (although, aside from smallNORB, it is only a bit worse). An intuitive explanation for this is that Halfway loss is more constrained than Softmax in terms of what it demands of distribution of points in $\phi(\mathbf{x})$. While these constraints are demonstrably favourable to generalisation in representationally rich models, they might be getting in the way of class separation objective in representationally limited models. In other words, Halfway loss may provide a better objective for classification, but an objective that is a bit harder to attain in models with limited transformation dynamics.

We also found that the cost-sensitive learning aspect of the Halfway loss was critical for its good performance. This is most interesting given that previously work by Zhou & Liu (2006) found cost-sensitive learning not to be useful for multi-class one-against-rest optimisation, albeit for different loss functions. The motivation for class weighting in Halfway loss is to ensure that the optimisation does not drive the variance around the negative margin to a smaller value than the variance around the positive margin. We take the need for class balancing in Halfway training as a confirmation of our assumption that an even reduction of variance around the positive and negative margin is critical for good generalisation.

The MSE-like nature of the Halfway loss has a disadvantage in that it presumes a normal distribution of the data around the margin. While it does succeed in minimising the margin variance, it also produces a symmetric distribution of data around the margin (as Figure 4b shows). It is possible that minimisation of variance while producing distributions skewed away from the margin might improve the generalisation even further.

7 Conclusion

We have taken the ideas around margin distribution from boosting theory and applied them to deep learning. The driving hypothesis of our work was that maximisation of the margin alone is not a useful objective for architectures where the feature extraction function changes during optimisation. However, minimisation of margin variance might be. We have provided some theoretical evidence that maximisation of margin in a neural network might be trivial. This we followed with empirical investigation of the importance of margin variance.

We proposed the Halfway loss function as the training objective that minimises the normalised margin variance . It's an MSE-like training objective with cost-sensitive learning that aims to reduce variance around halfway point between 0 and maximum margin value (as calculated from the training dataset). Our empirical evaluation on known image datasets demonstrates superiority of Halfway over the Softmax Cross-Entropy loss in representationally rich fully connected, as well as convolutional, neural networks. We also confirmed that in the balance of things, Halfway loss does seem to provide better generalisation - in terms of producing a validation test score that is a better estimation of the test score, while ensuring better test data performance.

For the future work, given the empirical evidence this work presents, we believe it would be worthwhile to find theoretical proofs that establish the significance of margin variance as well as the irrelevance of the margin mean for generalisation in deep architectures. On the empirical side, it might be also possible to form better loss functions which minimise margin variance but do not enforce symmetric distribution of the points around the margin, and thus possibly lead to even better generalisation.

Acknowledgements

We gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan X GPU used for this research.

References

Abadi, Martín, Agarwal, Ashish, Barham, Paul, Brevdo, Eugene, Chen, Zhifeng, Citro, Craig, Corrado, Greg S., Davis, Andy, Dean, Jeffrey, Devin, Matthieu, Ghemawat, Sanjay, Goodfellow, Ian, Harp, Andrew, Irving, Geoffrey, Isard, Michael, Jia, Yangqing, Jozefowicz, Rafal, Kaiser, Lukasz, Kudlur, Manjunath, Levenberg, Josh, Mané, Dan, Monga, Rajat, Moore, Sherry, Murray, Derek, Olah, Chris, Schuster, Mike, Shlens, Jonathon, Steiner, Benoit, Sutskever, Ilya, Talwar, Kunal, Tucker, Paul, Vanhoucke, Vincent, Vasudevan, Vijay, Viégas, Fernanda, Vinyals, Oriol, Warden, Pete, Wattenberg, Martin, Wicke, Martin, Yu, Yuan, and Zheng, Xiaoqiang. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. URL `http://tensorflow.org/`. Software available from tensorflow.org.

Breiman, Leo. Prediction games and arcing algorithms. *Neural Computation*, 11(7):1493–1517, 1999.

Cortes, Corinna and Vapnik, Vladimir N. Support vector networks. *Machine Learning*, 20(3): 273–297, 1995.

Cristianini, Nello and Shawe-Taylor, John. *An Introduction to Support Vector Machines : and other Kernel-based Learning Methods*. Cambridge University Press, Cambridge, NY, 2000.

Elkan, Charles. The foundations of cost-sensitive learning. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence*, volume 2 of *IJCAI'01*, pp. 973–978. Morgan Kaufmann Publishers Inc., 2001.

Gao, Wei and Zhou, Zhi-Hua. On the doubt about margin explanation of boosting. *Artificial Intelligence*, 203:1 – 18, 2013.

Glorot, Xavier, Bordes, Antoine, and Bengio, Yoshua. Deep sparse rectifier neural networks. In Gordon, Geoffrey J. and Dunson, David B. (eds.), *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics (AISTATS-11)*, volume 15, pp. 315–323. Journal of Machine Learning Research - Workshop and Conference Proceedings, 2011.

Hahnloser, Richard H. R., Seung, H. Sebastian, and Slotine, Jean-Jacques. Permitted and forbidden sets in symmetric threshold-linear networks. *Neural Computation*, 15(3):621–638, 2003.

Ioffe, Sergey and Szegedy, Christian. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *CoRR*, abs/1502.03167, 2015.

Jayadeva, Deb, Alok Kanti, and Chandra, Suresh. Binary classification by svm based tree type neural networks. In *Neural Networks, 2002. IJCNN '02. Proceedings of the 2002 International Joint Conference on*, volume 3, pp. 2773–2778, 2002.

Kingma, Diederik P. and Ba, Jimmy. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014.

Krizhevsky, Alex. Learning multiple layers of features from tiny images. Technical report, Computer Science Department, University of Toronto, 2009.

Lanckriet, Gert R. G., Cristianini, Nello, Bartlett, Peter, Ghaoui, Laurent El, and Jordan, Michael I. Learning the kernel matrix with semidefinite programming. *Journal of Machine Learning Research*, 5:27–72, 2004.

LeCun, Yann, Bottou, Léon, Bengio, Yoshua, and Haffner, Patrick. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86:2278–2324, 1998.

LeCun, Yann, Huang, Fu Jie, and Bottou, Léon. Learning methods for generic object recognition with invariance to pose and lighting. *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR) 2004*, 2:97–104, 2004.

Ludwig, Oswaldo and Nunes, Urbano. Novel maximum-margin training algorithms for supervised neural networks. *IEEE Transactions on Neural Networks*, 21(6):972–984, 2010.

Nishikawa, Takahiro and Abe, Shigeo. Maximizing margins of multilayer neural networks. In *Neural Information Processing, 2002. ICONIP '02. Proceedings of the 9th International Conference on*, volume 1, pp. 322–326, 2002.

Reyzin, Lev and Schapire, Robert E. How boosting the margin can also boost classifier complexity. In *Proceedings of the 23rd International Conference on Machine Learning*, ICML '06, pp. 753–760, New York, NY, USA, 2006. ACM.

Schapire, Robert E. The strength of weak learnability. *Machine Learning*, 5:197–227, 1990.

Schapire, Robert E., Freund, Yoav, Bartlett, Peter, and Lee, Wee Sun. Boosting the margin: a new explanation for the effectiveness of voting methods. *The Annals of Statistics*, 26(5):1651–1686, 1998.

Vapnik, Vladimir N. *The Nature of Statistical Learning Theory*. Springer, New York, 1995.

Wang, Liwei, Sugiyama, Masashi, Jing, Zhaoxiang, Yang, Cheng, Zhou, Zhi-Hua, and Feng, Jufu. A refined margin analysis for boosting algorithms via equilibrium margin. *J. Mach. Learn. Res.*, 12: 1835–1863, 2011.

Young, Steven and Downs, Tom. Carve-a constructive algorithm for real-valued examples. *IEEE Transactions on Neural Networks*, 9(6):1180–1190, 1998.

Zhang, Teng and Zhou, Zhi-Hua. Large margin distribution machine. *CoRR*, abs/1311.0989, 2013.

Zhang, Teng and Zhou, Zhi-Hua. Optimal margin distribution machine. *CoRR*, abs/1604.03348, 2016.

Zhou, Zhi-Hua and Liu, Xu-Ying. On multi-class cost-sensitive learning. In *Proceedings of the 21st National Conference on Artificial Intelligence*, volume 1 of *AAAI'06*, pp. 567–572. AAAI Press, 2006.